the network security company™

Re-Inventing Network Security to Safely Enable Applications

The Next-Generation Firewall Forms a New Foundation

November 2012

Executive Summary

It's no secret that modern applications and threats easily circumvent port-blocking firewalls, rendering ineffective what has long been the cornerstone of network security. Neither is it surprising that attempts to remedy this situation by deploying firewall "helpers," bolting application awareness onto existing firewall products, or turning to Unified Threat Management devices have all failed. The bottom line is that approaches which continue to classify traffic based only on ports and protocols are incapable of enabling the emerging generation of applications, users, and infrastructures.

To address this situation, Palo Alto Networks is re-inventing network security from the ground up. By focusing on applications, users and content—not ports and protocols—we're delivering a truly innovative platform that provides enterprises with the visibility and control required to safely enable modern applications, without having to absorb the risks that typically accompany them. At the core of this platform, the Palo Alto Networks Next-Generation Firewall also helps enterprises simplify their network security infrastructure.

Built on a high-performance architecture and incorporating a rich set of security, networking, and management functionality, our next-generation firewall significantly reduces network complexity and cost of ownership by eliminating the need to deploy a wide variety of supplemental network security products. A set of innovative technologies and enterprise-class capabilities further solidify enterprise gains by ensuring not only that consistent security and application enablement is available for all enterprise users and network locations, but also that the solution remains effective despite the onslaught of advanced malware and targeted attacks.

Network Security is No Longer Black and White

The old model for network security was simple because everything was black and white. Business applications constituted good, low-risk traffic that should be allowed, while everything else—including threats—constituted bad traffic that should be stopped. But that's no longer the case, primarily for three reasons.

Today's applications are mostly gray—not black and white. These days, personal productivity and lifestyle applications that enable people to handle their non-work affairs and maintain online personas run alongside corporate productivity applications used to automate key business processes. The problem is not so much the growing diversity of applications, but the inability to accurately classify any given application as good or bad – particularly when which end of the spectrum many of them fall on can vary from one user/session to the next.

For example, a personal productivity application used to share product documentation with a prospective customer would be "good" (medium risk, high reward), while using the same application to forward details of an upcoming release to a "friends list" that includes employees of a competitor would be "bad" (high risk, low reward).

A modern network security solution must be able not only to distinguish one type of application from the next, but also to account for other contextual variables that are applicable in any given scenario – such as who is using the application, from what type of device, and for what purpose.

Today's applications are evasive. Several other factors further complicate the core task of "distinguishing one type of application from the next." One challenge in particular is that in order to maximize their usability, many personal productivity and lifestyle applications have been designed to circumvent traditional firewalls by dynamically adjusting how they communicate.

Common tactics include:

- Port hopping, where ports/protocols are randomly shifted over the course of a session;
- Use of non-standard ports, such as running Yahoo! Messenger over TCP port 80 instead of TCP port 5050;
- Tunneling within commonly used services, such as when P2P file sharing or an IM client like Meebo runs over HTTP; and,
- Hiding within SSL encryption.

Another vexing issue is that many business applications are now being designed to take advantage of these same techniques. Examples include Box, Microsoft SharePoint and Salesforce.com. This is typically done to facilitate operation in the broadest set of scenarios and with the least amount of disruption for customers, partners, and the organization's own security and operations departments, but the unintended side effect for IT is further loss of control over network communications.

Even something as seemingly benign as the ongoing "webification" of enterprise applications is problematic, as legacy network security products are unable to resolve individual applications from the HTTP and HTTPS sessions that constitute over two thirds of all enterprise traffic.

Today's threats are coming along for the ride. The shift in motivation for today's hackers—from building reputations to actually making money – also brought applications squarely into their sights. After all, threats designed to operate at the application layer have multiple opportunities to circumvent an organization's defenses. Such threats not only pass right through legacy security infrastructure built only to provide network-layer protection, but can also obtain "free passage" into enterprise networks by targeting applications that utilize the aforementioned evasion techniques. In fact, even products that have been designed to accurately distinguish individual applications will be susceptible if they don't also incorporate threat prevention capabilities that enable them to reliably identify and stop threats riding along with applications the organization designates as "allowed."

IT is No Longer in Control

To be clear, the problem is not that there have been substantial changes to the application and threat landscapes; rather, it's that commonly deployed network security technologies have failed to keep pace.

Relying on Legacy Port-Blocking Firewalls is Like Flying Blind

Because they are deployed in-line at critical network junctions, firewalls see all traffic. This makes them the ideal resource to provide granular access control. It turns out, however, that most firewalls are far-sighted. They can see general shapes, but not the finer details of what's actually happening. This is because they operate by inferring the application-layer service that a given stream of traffic is associated with based on port numbers. They rely on a convention —not a requirement—that a given port corresponds to a given service (e.g., TCP port 80 corresponds to HTTP). This means they're also incapable of distinguishing between different applications that use the same port/service.

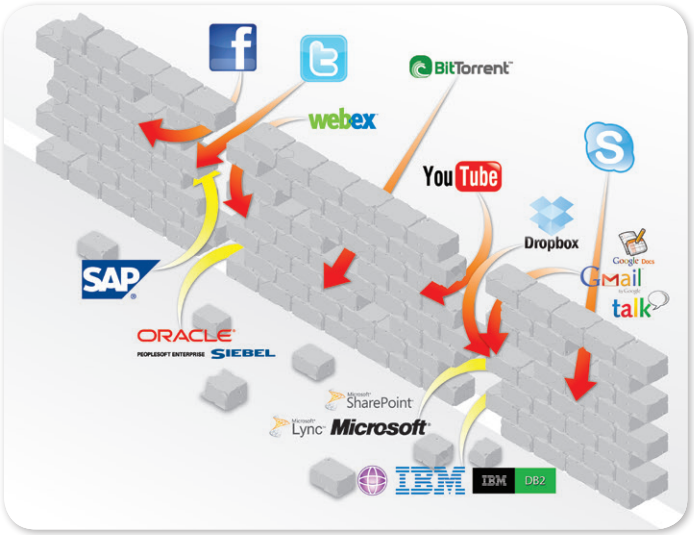

Figure 1: Port Blocking Firewalls Can't See or Control Applications

From a technological perspective, the impact of changes to the application landscape is that traditional, port-blocking firewalls have essentially gone blind. Not only are they unable to account for common evasion techniques, but they also lack the visibility, contextual awareness, and intelligence to discern:

- which network traffic corresponds to applications that serve a legitimate business purpose;
- which network traffic corresponds to applications that can serve a legitimate business purpose but, in a given instance, are being used for unsanctioned activities; and,
- which network traffic, even though it corresponds to legitimate business activities, should be blocked because it includes malware or other types of threats.

Moreover, their response options are too coarse-grained. They can either block or allow traffic, but offer little variation in between to adequately account for all of the "gray" applications enterprises ultimately need to support—for example, by allowing certain functions within an application but not others, applying traffic-shaping policies to certain traffic, or controlling access based on users, groups, or time of day.

From an operational perspective, the net result is that within many organizations IT has lost control. The inability of their security infrastructure to effectively distinguish good/desirable applications from those that are bad/unwanted has left them with an all or nothing decision. They can take a permissive stance that ensures the accessibility of important applications but also allows unwanted ones to proceed as well. Or they can just say "no"—an approach that maintains a high state of security but precludes the use of beneficial gray-area applications while increasing the risk of accidentally blocking essential business applications.

Firewall Remedies Have Failed

Unfortunately, none of the steps previously taken to address the inadequacies of traditional firewalls have given today's IT departments what they really need: the confidence to say "yes" to requests made by the business based on having the ability to exert granular control and provide in-depth protection down to the level of individual applications.

Bolting-on Deep Packet Inspection is Fundamentally Flawed

Attempting to correct the myopic nature of a traditional firewall by incorporating deep packet inspection (DPI) capabilities might appear to be a reasonable approach. However, this tactic only leads to an incremental boost in security effectiveness. The issue is that the new functionality is integrated rather than embedded, and the port-blocking firewall, with its complete lack of application awareness, is still used for initial classification of all traffic. As a result:

- Not everything that should be inspected necessarily gets inspected. Because the firewall is unable to accurately classify application traffic, deciding which sessions get passed to the DPI engine becomes a hit-or-miss proposition.
- Policy management gets convoluted. Rules on how to handle individual applications get "nested" within the DPI portion of the product—which itself is engaged as part of a higher/outer level access control policy.
- Inadequate performance forces compromises to be made. Inefficient use of system resources and CPU and memory intensive application-layer functionality put considerable strain on the underlying platform. To account for this situation, administrators can only implement advanced filtering capabilities selectively.

Deploying Firewall "Helpers" Only Creates Another Problem

Suggesting that enterprises compensate for their firewall's deficiencies by deploying a collection of additional, standalone security products—such as intrusion prevention, network AV, URL filtering, and data loss prevention appliances—has an eerily similar outcome, with one added twist.

To begin with, not everything that should get inspected does because the firewall helpers aren't exposed to all of the traffic, rely on the same port/protocol-based classification techniques that have failed the legacy firewall, or only provide coverage for a limited set of applications. Policy management is an even greater problem due to access control and inspection rules being spread across several consoles and involving multiple policy models. And performance is still an issue too, at least in terms of having a relatively high aggregate latency.

Then comes the twist: device sprawl. As one "solution" after another is added to the network, device count, complexity, and total cost of ownership all continue to rise. The result is an unwieldy, ineffective, and costly endeavor that is simply not sustainable.

UTM Only Makes What is Broken Cheaper

Another proposed remedy, the Unified Threat Management (UTM) device at least helps alleviate device sprawl. Instead of having a bunch of "helper" countermeasures deployed as separate devices, with UTM they all come in one physical package. But so what? The result is really no different than the bolted-on approach and, therefore, exhibits the same deficiencies. Inadequate application classification and blind spots in the inspections that are performed remain as fundamental problems, while performance and policy management issues are compounded based on having to account for multiple additional countermeasures instead of just one (i.e., DPI).

It's Time to Re-invent Network Security

The bottom line is that network security in most enterprises is fragmented and broken, exposing the business to unwanted risks and ever-rising costs. Traditional network security solutions have simply failed to keep pace with changing conditions, and the remedies put forth to compensate for their deficiencies have proven ineffective.

What today's enterprises need instead is a solution that:

- Enables IT to confidently say "yes" to whatever applications are needed to best support the business – by delivering the ability to accurately identify and granularly control applications while also preventing a broad array of threats;
- Achieves comprehensive coverage – by providing a consistent set of protection and enablement capabilities for all users and networks, regardless of their location; and,
- Simplifies security and networking infrastructure and operations – by obviating the need for numerous standalone products, each with its own policy construct, logs, and management console.

Introducing the Next-Generation Firewall

From the outset, Palo Alto Networks™ mission has been to restore the firewall as the cornerstone of enterprise network security infrastructure by "fixing the problem at its core." Starting with a blank slate, a world-class engineering team took an application-centric approach to traffic classification in order to enable full visibility and control of all types of applications running on enterprise networks – new-age and legacy ones alike. The result is the Palo Alto Networks next-generation firewall – the only solution that fully delivers on the essential functional requirements for a truly effective, modern firewall:

- Identify applications regardless of port, protocol, evasive tactics or SSL encryption
- Enable fine-grained visibility and policy control over application access and functionality
- Identify and control users regardless of IP address, location, or device
- Protect against known and unknown application-borne threats
- Support multi-gigabit, in-line deployments with negligible performance degradation

The key to this distinction is the combination of several innovative technologies, a high-performance architecture, and a robust set of additional enterprise-class capabilities that deliver a complete, end-to-end network security solution.

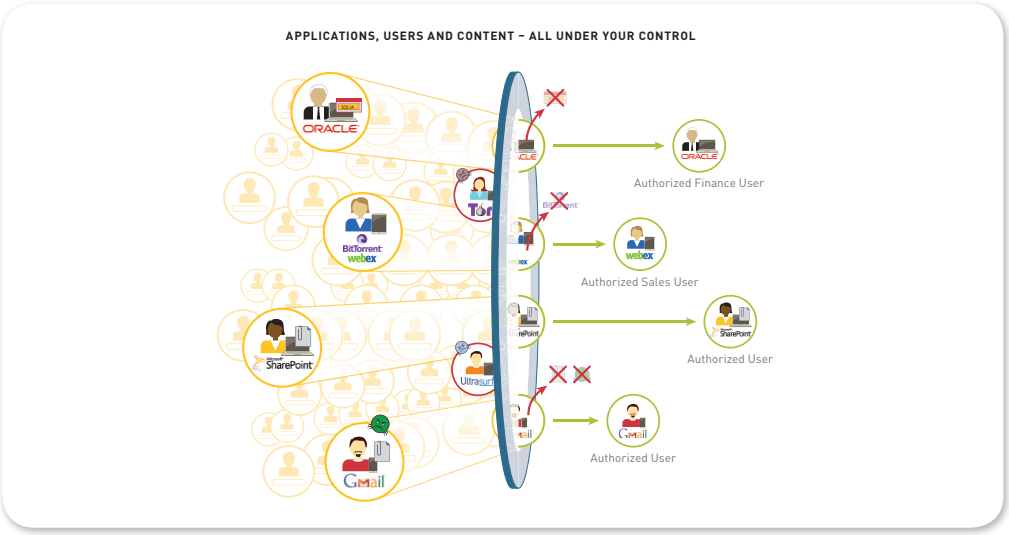

Figure 2: Next-Generation Firewalls Focus on Applications, Users, and Content

Unique Identification and Detection Technologies Restore Visibility and Control

A collection of unique identification and content detection technologies enable enterprises to focus on business relevant elements such as applications, users, and content for policy controls, instead of having to rely on nebulous and often misleading attributes such as ports and protocols. Also significant is the resulting unified management model, which dramatically simplifies security operations by providing administrators with one place to manage rules that seamlessly account for all relevant attributes and mitigations and one place to obtain and analyze all corresponding log information.

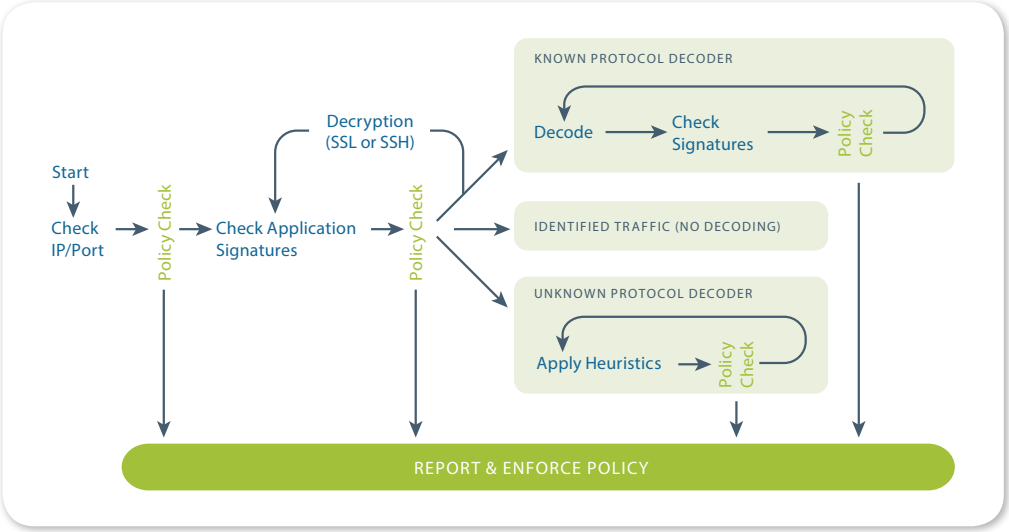

Figure 3: App-ID Identifies Applications Across All Ports

Positively Identify Applications

App-ID™ is the patent-pending traffic classification technology at the heart of the next-generation firewall. With App-ID, up to four distinct mechanisms are used to determine the exact identity of more than 1500 applications traversing the network—irrespective of port, protocol, evasive tactic or encryption.

Application Signatures. Used at multiple steps in the classification process, context-based signatures look for unique properties and transaction characteristics that identify applications regardless of the port or protocol being used.

SSL/SSH Decryption. If SSL or SSH encryption is in use and a decryption policy is in place, the traffic is decrypted and passed on to other mechanisms as needed.

Application Protocol Decoding. Decoders for known protocols validate conformance to the protocol specification, detect other applications tunneling inside of the protocol, and identify individual functions within a given application (e.g., WebEx Desktop Sharing).

Heuristics. Additional heuristic/behavioral techniques are engaged as needed to identify troublesome applications that typically elude advanced signature and protocol analysis, such as peer-to-peer or VoIP tools that use proprietary encryption.

Policy checks at multiple steps help guide the process and provide opportunities for granular control. In addition, a powerful application browser (ACC) delivers a wealth of intelligence about detected applications so that administrators can make well-informed decisions. With ACC, applications can be viewed by category, subcategory, underlying technology and characteristics such as: file transfer capabilities, the ability to evade detection, and the propensity to consume bandwidth, transmit malware, or otherwise be misused. For any "unknown" application traffic that is detected, administrators have the option of submitting a sample to Palo Alto Networks for development of a new App-ID, or generating their own custom App-ID.

With App-ID, IT obtains not only a clearer picture of what's happening on their networks, but also the ability to create and enforce policies that deliver secure application enablement.

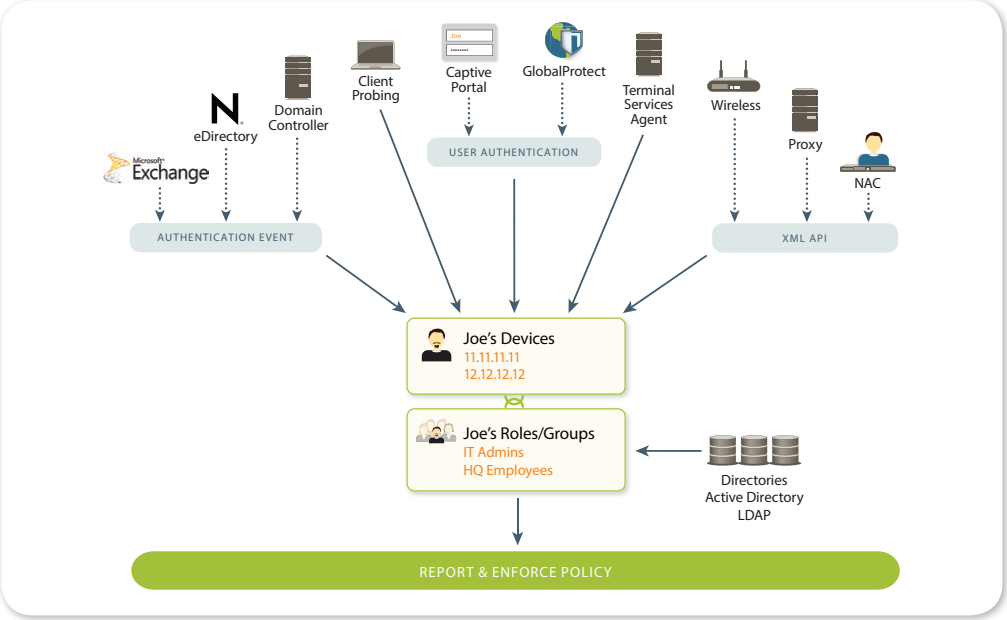

Figure 4: User-ID Integrates Enterprise Directories for User-based Policies, Reporting and Forensics

Securely Enabling Applications Based on Users and Groups

User-ID™ seamlessly integrates Palo Alto Networks firewalls with a range of enterprise directory and terminal services products, enabling administrators to tie application activity and security policies to users and groups—not just IP addresses. Working together, network-based User-ID Agents and next-generation firewalls use a combination of techniques—including login monitoring, end-station polling, standard LDAP queries, captive portal, and XML API-based interrogation—to maintain user-to-IP address relationships and harvest relevant details (e.g., group assignments). These details are then available to:

- Identify the specific users responsible for all application, content, and threat traffic on the network;
- Extend user-based application enablement policies across Microsoft Windows, MAC OS X, Apple iOS, Unix, and terminal services users; and,
- Facilitate user-based analysis, reporting and forensics.

With User-ID, IT departments get another powerful mechanism to enable intelligent control of applications. For example, a social networking application that would normally be blocked because of its risky nature can now be enabled for individuals or groups that have a legitimate need to use it, such as the human resources department.

High-performance Content Control and Threat Prevention

Like their counterpart technologies, Content-ID™ and WildFire™ infuse the Palo Alto Networks next-generation firewall with capabilities previously unheard of in an enterprise firewall. The net result is the ability to further control access through the firewall based on the nature of the content being communicated, and whether or not it contains any threats.

Intrusion Prevention and Network Antivirus. An embedded, NSS Labs Recommended intrusion prevention system protects networks from all types of vulnerability exploits, buffer overflows, DoS attacks and port scans. Traffic normalization eliminates invalid packets, while TCP reassembly and IP de-fragmentation counteract packet-level evasion techniques to ensure the utmost accuracy and protection. In addition, known viruses and malware are eliminated from allowed traffic using our own database of millions of malware samples and a stream-based scanning technique that minimizes latency by not waiting for entire files to be loaded into memory before starting the scanning process. Performance is further enhanced by leveraging a uniform threat signature format that makes it possible to detect both intrusions and viruses with a single pass of a single scanning engine.

Modern Malware Defense. Complementing the Content-ID threat prevention capabilities, WildFire delivers essential protection against the rising tide of custom, targeted, and otherwise unknown malware bombarding today's enterprises. Subject to administrator-configured policies, unknown files and correlating information detected by the firewall are submitted to the WildFire virtualized sandbox. A cloud-based component that eliminates the need for additional dedicated hardware at every ingress point to an enterprise's network, the sandbox provides virtual targets where potentially dangerous files can be executed and observed for malicious behavior. For samples identified as malware, new protection signatures are automatically generated, tested for accuracy, and then distributed to all Palo Alto Networks customers as part of the daily malware signature updates. Customers with a WildFire subscription receive priority delivery in less than one hour. They also gain access to a wealth of actionable intelligence about newly detected malware through the WildFire portal. Detailed reports document observed behaviors and provide useful information about targeted users, applications, URLs, and so forth.

All of the threat prevention capabilities discussed fully leverage App-ID, thereby ensuring that threats and unknown files are detected even when evasive techniques are employed.

URL Filtering. A fully integrated, on-box URL database with 20 million entries across 76 categories allows administrators to control web surfing and access to high risk sites. Support for custom entries and categories and the ability to pull additional entries into cache from a hosted, 180 million URL database help ensure nothing slips through the cracks.

File and Data Filtering. Administrators can implement policies that reduce the risks from unwanted exposure of confidential information by taking advantage of embedded capabilities to: block files by their actual type (i.e., not based on just their extension), filter traffic for sensitive data patterns such as credit card and social security numbers, and control the file transfer functionality within individual applications, such as instant messaging clients.

The bottom line is that with Content-ID and WildFire, IT departments get additional, powerful mechanisms for controlling network access along with fully integrated, broad-spectrum threat prevention – all without having to invest in a collection of standalone products and their separate policies, operations, and views of the world.

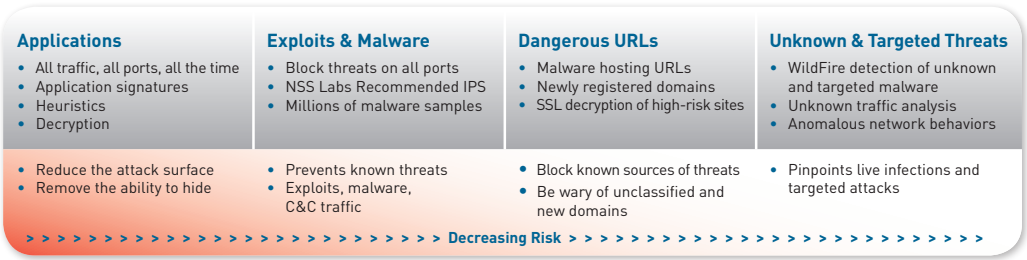

Figure 5: Next-Generation Firewall Delivers Fully Integrated, Broad-Spectrum Threat Prevention

High-Performance Architecture Delivers "Security Without Compromises"

Having a comprehensive suite of application awareness and content inspection capabilities makes little difference if administrators are unable to fully engage them due to performance constraints. Accordingly, Palo Alto Networks designed the next-generation firewall to have a single pass, parallel processing (SP3) architecture.

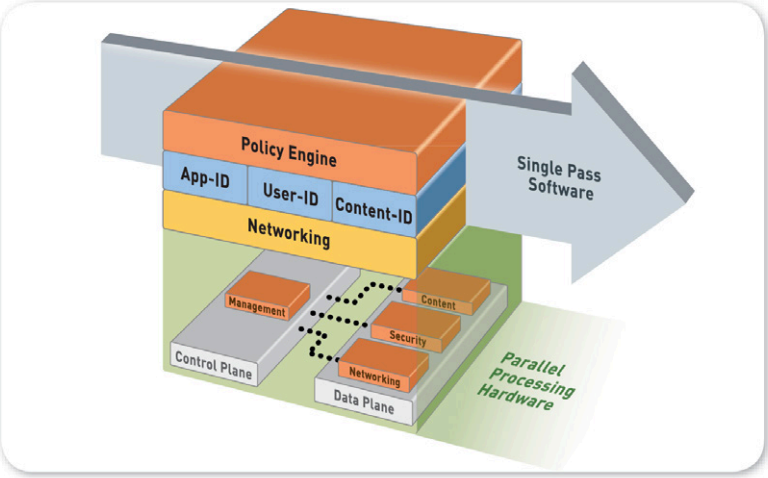

Figure 6: Single Pass Parallel Processing Architecture Optimally Combines Software and Hardware for Superior Performance

With most conventional network security products, each high-level function processes the traffic stream independently. The resulting multi-pass approach requires low-level packet handling and stream reassembly routines to be repeated numerous times. System resources are used inefficiently and considerable latency is introduced. In contrast, the processing model for the Palo Alto Networks next-generation firewall is highly structured and essentially linear. This eliminates repetitive handling of packets and streams, dramatically reducing the burden placed on system hardware and minimizing latency.

Performance is further enhanced by having separate control and data planes, with the latter having function-specific, parallel processing hardware engines, including:

- a network processor for initial packet handling and other network-layer functions;
- a multi-core security processor and hardware-based acceleration for standardized functions; and,
- a dedicated content scanning engine.

Delivering up to 20 Gbps of low-latency, App-ID throughput, the net result is a solution that delivers next-generation visibility and control of applications, users, and content without having to make any compromises.

Fixing the Firewall is Only the First Step

Although it is an essential starting point, fixing the firewall is only the first step toward delivering truly effective network security. After all, it's not just the application landscape that has changed in recent years. Today's IT managers are also struggling to account for:

- an increasingly mobile workforce;
- the growing adoption of public cloud services; and
- the transition to highly virtualized, cloud-like data centers.

All of these trends are having an impact not only on what capabilities are needed to safely enable applications, but, to an even greater extent, on where those capabilities are needed. Simply put, it's no longer enough to implement firewalls solely at the network perimeter. Accordingly, another mission for Palo Alto Networks has been to ensure that our next-generation firewall is applicable for all users and network locations requiring robust network security and safe application enablement.

Enterprise-class Capabilities Ensure Consistent Security Everywhere

The days of relying primarily on perimeter-based defenses for network security are all but gone. A host of factors—including geographically distributed offices, mobile users, and even local users with mobile devices—now dictate the need for a far greater number of security enforcement points. The goal of Palo Alto Networks in this regard is to deliver a solution that accounts for the unique requirements of all users and locations yet still provides a consistent set of protection and application enablement capabilities—all without having to manage a completely independent set of policies and infrastructure.

Specific ways Palo Alto Networks provides coverage for more than just the network perimeter include the following:

Palo Alto Networks for the Data Center. Suitability for deployment in enterprise data centers begins with the SP3 Architecture and a core design that delivers robust network security without having to make compromises due to performance considerations. Equally important is the ability to stop modern threats and rogue applications from operating in an organization's "inner sanctum." Support for a broad range of networking technologies (e.g., L2/L3 switching, dynamic routing, 802.1Q VLANs, trunked ports, traffic shaping, etc.), deployment modes, and high availability configurations also ensures the ability to "fit in" to whatever data center design/architecture an enterprise might have. Additional data center-friendly capabilities include the following:

- Support for segmentation and multi-tenant configurations. A virtual systems capability allows administrators to configure multiple, independent firewall instances within a single physical firewall appliance. Each instance has its own set of policies and management interface, providing a convenient and economical approach for IT to segment applications, servers, and data resources (e.g., for compliance purposes), or provide separate support for different business units and internal constituencies.

- Support for enterprise cloud networks. With the introduction of the Palo Alto Networks VM-Series, enterprises now have the option of deploying virtual machine-based instances of our next-generation firewall that feature the same application, user, and content control capabilities as with our physical appliances. A unified policy and management model that spans both VM-based and physical firewall appliances and a dynamic address object feature—which pins policies to migrating VMs/workloads—helps ensure consistent security while simplifying administration. In addition, an XML API delivers out-of-the-box integration with third-party orchestration and automation tools to enable dynamic firewall provisioning and management. The net result is a highly affordable and easy-to-implement option that enables IT managers to: (a) extend next-generation firewall coverage to locations and use cases where cost was previously an issue, (b) provide robust security for network-in-a-box configurations, (c) ensure policies and protections are maintained for migratory workloads, and (d) fully support the transformation to dynamic, cloud-like data centers.

Palo Alto Networks for Branch Offices. Enterprises can establish a consistent level of protection and secure application enablement across all of their offices and facilities, regardless of size, by taking advantage of a portfolio of a dozen firewall appliances that blanket the performance/throughput spectrum, from under 100 Mbps to 20 Gbps, and offer a wide variety of interface choices. This includes the newest addition to the family, the PA-3000 series, which provides even more price/performance options for intermediate-size sites. In addition to supporting direct-to-Internet activities for distributed users, new PAN-OS capabilities also enable simple, rapid configuration of large-scale site-to-site VPNs for secure inter-office communication.

Palo Alto Networks for Mobile Users. A persistent client component that can be installed on-demand, GlobalProtect solves the security challenges introduced by roaming users by extending the same next-generation firewall policies and protection enforced within the physical perimeter to all users, no matter where they are located. When a remote user logs into their device, Global-Protect automatically determines the closest next-generation firewall, transparently authenticates the user, and establishes a secure connection (using SSL or IPSec). Users can then access corporate network and Internet-based resources subject to the same policies and protection that apply when operating locally. A consistent user experience and level of security is provided without the need to create and manage multiple, separate sets of policies. The only difference is that policies can expanded to also account for the security status of the user's device, such as whether the latest patches are installed.

Tying everything together is Panorama, a centralized management system that can optionally be deployed on available server hardware or Palo Alto Networks M-100 management appliances. Sharing the same look and feel as the device-level command line and web interface management options, Panorama simplifies network security management by providing a single, unified console for conducting all policy management, software and content update, logging, and reporting tasks across all components of the Palo Alto Networks solution.

A New Foundation for Securing Networks and Safely Enabling Applications

The Palo Alto Networks next-generation firewall provides today's enterprises with precisely what they need to take back control of their networks, to stop making compromises when it comes to information security, to put an end to costly appliance sprawl, and to get back to the business of making money. By delivering the visibility, control, and fully integrated threat protection required to safely enable applications for all users and network locations, Palo Alto Networks' innovative platform substantially raises the bar for security and operational efficiency while establishing a new foundation for enterprise network security.

paloalto NETWORKS
the network security company™

3300 Olcott Street
Santa Clara, CA 95054

Main: +1.408.573.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com